Technical Note 497

OVERLOADING INTENTIONS FOR EFFICIENT PRACTICAL REASONING

SRI Project 7363

Prepared by:

MARTHA POLLACK
Computer Scientist

Artificial Intelligence Center
Computing and Engineering Sciences Division

APPROVED FOR PUBLIC RELEASE: DISTRIBUTION UNLIMITED

This paper was supported by the Office of Naval Research Center under Contract No. N00014-89-C-005, by a contract with the Nippon Telegraph and Telephone Coporation and by a gift from the System Development Foundation

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **1990**	2. REPORT TYPE	3. DATES COVERED **00-00-1990 to 00-00-1990**
4. TITLE AND SUBTITLE **Overloading Intentions for Efficient Practical Reasoning**		5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **SRI International,333 Ravenswood Avenue,Menlo Park,CA,94025**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**		**35**	

Overloading Intentions for Efficient Practical Reasoning

Martha E. Pollack
Artificial Intelligence Center *and*
Center for the Study of Language and Information
SRI International
Menlo Park, CA 94025

1 Introduction

Agents, whether biological or artificial, have bounded reasoning capabilities. As a result, they cannot make reasoned decisions instantaneously; reasoning takes time. Yet all biological agents, and most artificial ones, are situated in dynamic environments, in which events not under their control occur. Agents in dynamic environments face a potential difficulty when they must make decisions about what to do—what goals to pursue and how to pursue them. They run the risk that the world may change while they are engaged in such practical reasoning; in fact, it may change in ways that undermine the very assumptions upon which their reasoning is proceeding. Agents who blindly push forward with every practical-reasoning problem they encounter, without regard to the amount of time they are taking or the changes meanwhile going on, are not proceeding in an intelligent fashion.

Dynamic environments and computational resource bounds thus pose a challenge that has led some researchers in Artificial Intelligence (AI) to propose that artificial agents be designed to avoid execution-time practical reasoning (Agre and Chapman 1987, Brooks 1991, Kaelbling 1988). To obviate the need for such reasoning, the designer of an agent must specify a function $\mathcal{F}$, from each set of situational features that the agent can directly observe, to appropriate actions it should perform. In addition, the designer should guarantee that each situation in which the agent might find itself corresponds to exactly one set of situational features in the domain of $\mathcal{F}$, so that there is a function from situations to actions.[1] Although the ultimate feasibility of this approach remains an open question, many researchers are skeptical about the practicality of relying exclusively on such design-time decision making for agents who are to be situated in complex environments (D'Ambrosio and Fehling 1989, Doyle 1988, Pollock 1989).

An alternative is to allow execution-time reasoning, but to develop techniques for controlling it. One way to provide such control is to enable the agent to perform explicit meta-level reasoning—reasoning about when to reason and when to act instead. Another approach is to include in the agent's architecture mechanisms that directly control the reasoning it performs. In an earlier paper (Bratman *et al.* 1988), Bratman, Israel, and I took this latter approach, specifying an architecture for practical reasoning that built upon Bratman's (1987) theory of the role of an agent's plans in constraining subsequent practical reasoning. Specifically, Bratman suggested two ways in which an agent's plans have this constraining effect:

- An agent's plans focus subsequent means-end reasoning—reasoning about how to achieve certain goals. Agents can, in general, concentrate on elaborating their existing plans, rather than on computing means to any desired ends.

- An agent's plans restrict the set of potential courses of action that need to be seriously considered. In general, agents can filter out options that are inconsistent with the performance of what they already plan to do.

In this paper, I argue that there is an additional way in which an agent's plans can be used to constrain practical reasoning: they can suggest solutions to means-end reasoning problems that the agent subsequently encounters. Moreover, such solutions can often be accepted without further deliberation about possible alternatives. An agent will often be able to guide its search for a way to achieve some goal G by looking for an action A that it already intends that can also subserve G, or, as I shall say, by looking for an intention that can be *overloaded*.[2] If it is successful in this, it can typically avoid attempting to find alternative ways of achieving G; it need not weigh the solution involving A against competing options. I will argue that such a strategy, fine-tuned in appropriate ways, is rational, despite the fact that it may sometimes lead to suboptimal behavior.

2 A Strategy for Means-End Reasoning

We begin with an example. I plan to go swimming at lunchtime today with my colleague Nan, and to stop on the way back to my office to buy a sandwich. Each of these plans includes several steps: for instance, as part of my plan to go swimming with Nan, I intend to meet her, walk over to the pool with her, and so on. As lunchtime approaches, I discover that I forgot to go to the bank before work, and consequently have no money with me. I thus face a practical reasoning problem: I must determine a way to get some money to pay for my sandwich. My existing plans suggest a way: when I meet Nan to walk to the pool, I can borrow money from her. Assuming I have reasonable confidence in the likely success of such a plan, I need not generate and consider alternative ways to get the money, such as stopping at the automatic teller machine, or borrowing money from John, or from Michelle, or from Deborah, or so on. Instead, I directly adopt the plan to borrow money from Nan when I meet her. Before discovering that I had no money, I intended to meet Nan for one purpose—it was part of my plan to go swimming with her. But now, that intention is overloaded, subserving two distinct goals: it is also part of my plan to get money for lunch.

To understand the reasoning illustrated by this example, we need to distinguish between two kinds of practical-reasoning processes, both of which I mentioned earlier: reasoning about what goals to pursue and reasoning about how to pursue them. The former, which is often called deliberation, has traditionally been the focus of decision theory (Savage 1972, Jeffrey 1983). The latter, means-end reasoning, has, until quite recently, been the primary focus of most research in AI planning theory. Indeed, in major AI planning systems such as STRIPS (Fikes and Nilsson 1971), NOAH (Sacerdoti 1977), SIPE (Wilkins 1988), and NONLIN (Tate 1977), the task performed is precisely that of finding some sequence of actions, that, when performed starting from a given state, will achieve a specified goal. Such a sequence of actions is called a plan for the goal. The term "plan," used in this way, does not entail an agent's commitment to performing, or even attempting to perform, the actions that constitute the plan. Rather, on this view, plans can be seen as "recipes" for achieving a goal. The decision to adopt, or form an intention to perform, some plan is not taken to be a necessary consequence of means-end reasoning.[3] I shall use the term "plan" in this way, contrasting it with adopted plans, those the agent actually intends.

AI models of means-end reasoning have typically cast it as a search process, in which nodes in the search space correspond to plans to achieve a given goal. These plans may be only partially specified. The search space is expanded by elaborating, or further specifying, one of the plans already contained in it (Georgeff 1987). There are various ways in which a plan can be elaborated. For instance, a plan that includes an action that is not directly executable can be elaborated by specifying a particular way of carrying out that action; a plan that includes a set of actions can be elaborated by imposing a temporal order on the members of the set; and a plan that includes an action involving objects whose identities are so far underspecified can be elaborated by fixing the identities of one or more of the objects. In performing means-end reasoning, a system will select some node specifying a partial plan, elaborate it in some way to form a new node, and then repeat the entire process. The process terminates when an elaboration results in a completely specified, and thus, executable, plan.[4]

AI planning systems thus make use of the fact that plans are hierarchically structured: elaboration can be viewed as the process of embedding a new, more specific plan within an existing one. As Bratman (1987) notes, "[p]lans concerning ends embed plans concerning means and preliminary steps; and more general intentions ...embed more specific ones" (p. 29). The distinction between these two kinds of embedding recurs in the AI literature. For instance, Kautz (1990) identifies two relations: (1) *decomposition*, which relates a plan to another plan that constitutes a way of carrying it out (means and preliminary steps), and (2) *abstraction*, which relates a specific plan to a more general one that subsumes it. It is useful to have a term to refer to the inverse relation to abstraction: I shall speak of this as *specialization*.

It is surprisingly difficult to give a precise statement of the difference between decomposition and specialization, despite the fact that, in many cases, the difference is intuitively clear. Kautz, Bratman, and others distinguish between them primarily by example. Kautz observes that a plan to make a pasta dish may be decomposed into the steps of making noodles, cooking noodles, and making sauce; and that it may be the abstraction of a plan to make spaghetti marinara. (And, thus, a plan to make a pasta dish may be specialized into one to make spaghetti marinara.) Bratman gives the example of a general plan to attend a concert being specialized into one of attending the Alma Trio's concert tonight. Within the framework of traditional AI models of planning, actions are represented with operators that encode the change in state that will occur when the action is performed. Specialization corresponds to the binding of one or more of an operator's parameters, or to the imposition of temporal constraints on a set of operators. Decomposition corresponds to operator expansion, that is, the transformation of a single operator into a set of more primitive operators that constitute a plan for it. Of course, this distinction hinges on the way in which the operators used to model a given domain are defined by the system builder. In any case, it is not necessary, for our purposes, to have a precise account of the difference between decomposition and specialization, as I will argue that overloading is a mechanism that applies to both kinds of plan elaboration.

Let us return to the example given earlier. There I faced a particular means-end reasoning problem: I needed to determine a way to get the money. Using the framework for means-end reasoning just described, my task is to elaborate a partial plan so far consisting only of the action of getting some money for lunch. What the example suggests is that to perform this elaboration, I first consider the actions I already intend, to see whether any of them can be part of an elaboration. In the example, I succeed, determining that my intended action of meeting Nan can also be used in an elaboration of my plan to get money. The elaboration step actually involves the introduction of two hierarchical levels: I decompose my plan to get money into a plan to borrow money from Nan, and then decompose this into a sequence of steps that begins with my meeting her.

It is useful to diagram informally the intentions an agent has at various stages of her reasoning process. In the current example, assuming that I go ahead and adopt the plan to get money by borrowing it from Nan, we can sketch the relevant portion of my intentions as follows:

go swimming with Nan
↑
{*meet Nan*, walk to swimming pool, . . .}

get some money for lunch
↑
borrow money from Nan
↑
{*meet Nan*, ask to borrow money from her, be given money}

In this diagram, the labels refer to the actions that are components of my adopted, that is, actually intended, plans, and the arrows represent elaboration relations amongst the actions. The action described at the head of an arrow is elaborated by the action(s) described at its tail. Or, to put it in more natural terms, I intend to achieve the action described at the head of the arrow by performing those described at the tail. Italicized labels are used to indicate the overloaded action: both mentions of meeting Nan refer to the same action.

In this example, the attempt to overload led to a particular decomposition of my new goal. But overloading may also suggest particular specializations of a new goal. Imagine a slight variation of the example, in which I first decide that the way I will obtain money is to borrow it from someone. (Perhaps I often forget my lunch money, and so have just gotten into the habit of borrowing from one colleague or another; hence, this is an easy—almost automatic—reasoning step for me to perform.) At this point in my reasoning, my adopted plans include the following:

go swimming with Nan
↑
{meet Nan, walk to swimming pool, . . .}

get some money for lunch
↑
borrow money [from someone]

Of course, my means-end reasoning task is still not complete, since, before I can execute my plan to borrow money, I need to determine from whom to borrow it. At this point, I may consider my existing plans, and recognize that I can overload my intended action of meeting Nan. In this case, overloading suggests a specialization of my intended action of borrowing money from someone, namely, borrowing it from Nan. With this specialization, I can overload my action of meeting Nan, using it both in the decomposition of my plan to go swimming with her, and in a decomposition of my plan to borrow money from her:

go swimming with Nan
|
{*meet Nan*, walk to swimming pool,...}

get some money for lunch
↑
borrow money [from someone]
↑
borrow money from Nan
↑
{*meet Nan*, ask to borrow money from her, be given money}

So overloading may suggest a particular way of elaborating a plan—and the elaboration may involve either decomposition or specialization. In addition, overloading may sometimes suggest a way of elaborating two (or more) plans at once. Consider the following example. I plan to clean out my garage this weekend. This entails disposing of various objects, including empty cardboard boxes, an unused but working bicycle, and an old, broken stereo system. I subsequently learn that my daughter's daycare center is having a yard sale, to which the parents are being asked to contribute. My existing intention to dispose of the bicycle, which is part of my larger plan to clean my garage, suggests a way to contribute to the yard sale: I can dispose of the bicycle by donating it to the sale.

Again, it is useful to sketch the plans as they are elaborated during the reasoning process described in the example. At the beginning of the means-end reasoning process, my adopted plans include the following:

clean garage donate [something] to yard sale

{dispose of boxes, dispose of bike,...}

There is no action that I already intend that can also be used in an elaboration of my new goal of contributing to the yard sale. However, I recognize that, because I'm going to dispose of my bike anyway, I might as well donate it to the sale. In other words, a simple elaboration of my existing plan to clean the garage yields a possibility for overloading. If I decompose my intention to dispose of the bike into an intention to donate it to the yard sale, I can also use the latter as a specialization of my new goal of donating something to the yard sale. My resulting plans look like this:

clean garage donate [something] to yard sale

{dispose of boxes, dispose of bike,...} *donate bike to yard sale*

donate bike to yard sale

What we have seen is that overloading can serve as a strategy for focusing means-end reasoning. When an agent is faced with a means-end reasoning problem—when she needs to elaborate a partial plan—she can begin by considering actions she already intends for other purposes, looking to see whether any of them can be used in solutions to the current problem. If this fails, she may attempt to generate some elaborations of existing plans, to see whether any of these result in candidates for mutual elaboration. But is it advantageous to reason in this way? Before answering this question, we need to turn to the second practical-reasoning process I mentioned earlier, deliberation.

3 Bypassing Deliberation

Deliberation is the process of deciding among competing alternative courses of action. One important class of deliberation problems involves deciding amongst alternative ways to achieve a goal, or, in terms of the model sketched in the preceding section, deciding amongst alternative elaborations of some plan.

In (Bratman *et al.* 1988), we proposed a two-stage deliberation architecture, based directly on Bratman's (1987) observation that agents are committed to doing what they plan, and that, as a result, they need not engage in full-blown deliberation about all alternatives that may be presented to them. Instead, suggested alternatives are first subjected to a filtering process, which generally screens out options that are incompatible with the agent's existing plans. (The filtering can be overridden by a mechanism that detects options that *prima facie* have some special significance.) A second stage then processes the surviving alternatives, weighing their relative strengths to select amongst them.

Certainly any candidate plan that is suggested as the result of overloading must be subject to filtering, that is, it must be determined to be compatible with the agent's already adopted plans. It will not do for me to form the intention to borrow money from Nan if I have a standing intention never to borrow money from anyone. However, a plan P that was suggested by overloading is more likely to be compatible with the agent's adopted plans than is a random elaboration of the new goal, because at least one part of P, the overloaded action, is already intended by the agent and thus has already been deemed in at least one context to be compatible. Of course, there may be prohibitions against the action being used in the new context, and there may also be prohibitions against other parts of the plan; the latter is illustrated by the case in which I intend never to borrow money.

Consider, however, cases in which P, the plan suggested by overloading, is, in fact, compatible with the agent's adopted plans. In these situations, the agent may well be justified in adopting P without computing a larger set of alternatives, the members of which she would then need to weigh against one another. Why? Because it is often a good idea to "kill two birds with one stone," to put one's actions to more than a single use. By performing one action in service of more than one goal, the agent will, in general, act efficiently. Although P may not turn out to be the best plan that the agent could use to achieve her goal, it is likely to be a reasonably good plan, given the efficiency of action inherent in it. And if she directly adopts P, the agent achieves a savings in reasoning cost: she avoids both computing alternatives to P and performing the deliberation needed to choose amongst P and its alternatives. In addition, adoption of P may reduce the amount of subsequent means-end reasoning needed: the agent will need to determine how to do only one action—the overloaded one that subserves two or more goals—rather than how to do several actions, one for each goal.

So a plan containing an overloaded action will exhibit a certain efficiency of action, and directly adopting it will result in an efficiency of reasoning. At least in everyday practical-reasoning tasks, in which efficiency of action tends to be a significant factor in a plan's utility, it is likely that the savings in reasoning cost will outweigh the advantages of the "best" possible plan relative to the one suggested by overloading.

Thus, when an agent successfully uses overloading to solve a means-end reasoning problem, the resulting plan may be directly adopted, assuming that it is consistent with the agent's already adopted plans. This contrasts with the traditional decision-theoretic model in which a complete set of alternatives is first generated and then weighed against one another. As such, the model of planning using overloading recalls Harman's (1986) hypothesis that

> ...many of one's decisions are of necessity what we might call *simple decisions*. These arise when one finds oneself with a salient end E and one recognizes a salient means M that will get one E. In a simple case, one does not consider whether there might be some other means to E or some other end distinct from E that one might now obtain, and one disregards any other consequences of one's act. One simply forms the intention of getting E by doing M (p. 106).

Overloading provides a computational account of how, in some cases, an agent can recognize a salient means M that will get her E. A plan for, or means to, E that includes an action already intended for some other purpose is one that will typically be regarded as salient. Borrowing money from Nan is an intuitively salient solution to my problem—after all, I'm going to be meeting her to go swimming anyway! Salience is a notoriously difficult property to model computationally, but the focusing inherent in overloading captures a special case of it.

In addition, the account of overloading given here can contribute to an explanation of why simple decision making may often be good decision making. If the salient means M that the agent adopts is one that is suggested by overloading, then there is a good chance that M is a means that would be ranked relatively highly by a deliberation process. Thus, given an agent's resource bounds, it is reasonable for her to adopt a plan to do E by means of M, without attempting to generate and weigh other options. Of course, not all simple decisions will be instances of overloading, and we therefore need, in the future, to identify additional strategies for means-end reasoning that can explain why simple decision making works as well as it does.

A pair of empirical studies of human planning behavior show that humans often make use of something like overloading in performing practical reasoning. The first study, conducted by Hayes-Roth and Hayes-Roth (1979), investigated the behavior of human subjects who were assigned a complex practical-reasoning task involving the coordination of a number of errands, such as buying a fan belt, picking up medicine for one's dog, and meeting a friend for lunch; these were to be accomplished in a hypothetical town for which the subject was given a map. One of the primary observations made in this study was that subjects tended to plan opportunistically: because it was too difficult for them to compute an optimal schedule for all the errands, they tended, instead, to select a few of the most important tasks, schedule those, and then further develop their schedules around that framework. So, for example, a subject might determine that it was crucial to pick up the medicine for his dog, and thus decide on a definite time to do this. This decision would then influence a host of others. He might next decide to carry out all those tasks that could be performed at locations near the veterinary supply store, effectively overloading his trip to the area where the supply store is located. He might, for instance, decide to buy the fan belt on his way to get the medicine, since his route will take him past an automotive parts store. In fact, he might even select that particular automotive parts store, as opposed to the other ones in town, simply because he'll be passing by it.

More recently, Hayes (1989) observed the behavior of machinists' developing plans for machining particular parts. A plan to machine a part specifies, amongst other things, a sequence of set-ups, which describe how the part is to be oriented and clamped, and, for each set-up, a sequence of cutting steps, which describe the tool and method used. Hayes found that as machinists gain more expertise, they have an increasing tendency to construct plans that involve what she calls *operator overlap*.[5] Expert machinists often construct plans in which a single set-up is used for several cutting steps, even if some of the cutting steps then are uncommon ones. Novice machinists, in contrast, tend only to employ common cutting steps, with the result that their plans frequently involve multiple set-ups.

The difference between expert and novice behavior in this setting is quite interesting. Hayes determined that the plans produced by the experts indeed exhibited an efficiency of action: they took less time to perform than did the plans generated by novices. However, as Hayes notes, because the plans that include operator overlap make use of less common, and thus, presumably, less salient decompositions, they may be more difficult to generate. If this is so, then the extra means-end reasoning needed to compute the plans with operator overlap could, in principle, surpass the amount of time saved during performance. In other words, the efficiency in action that might be gained by overloading may be overridden by the inefficiency in means-end reasoning needed to use overloading in the first place.

It appears that, for the novices, it is relatively easy to generate the plans that comprise only frequently used actions. In analyzing protocols of the machinists describing their planning process, Hayes determined that the novices do not appear to perform extensive searches for optimal plans. Rather, they generate a single plan, performing the kind of simple decision making suggested by Harman. Overloading would thus not provide to them a savings in deliberation cost, as they are not, in any case, engaging in deliberation. Moreover, the potential savings in subsequent means-end reasoning may be undercut by the inclusion of uncommon steps in the plan: it may actually be easier for the novices to determine how to carry out two distinct actions that are each quite common than to determine how to carry out a single action of an infrequently used type. So overloading may provide little or no efficiency of reasoning for the novices. And, as I have already noted, the efficiency of action that it can provide may be outweighed by the additional means-end reasoning required initially to find the possibility of overloading.

But what of the experts? How is it that they manage to generate overloaded plans, without the burden of additional means-end reasoning becoming too great? Here again, the protocol analysis is revealing. Like the novices, the experts do not generate a large number of candidate plans; often they too generate only a single plan consisting of a sequence of frequently used subplans. However, the experts were able to detect certain cues in some parts specifications that indicated that overlap might be feasible. The fact that two angles were to be cut 90 degrees apart is an example of such a cue; it suggests that side-milling on a sine table may be appropriate. What the experts appear to do is to look for cues of this sort, which suggest to them plans involving operator overlap. The cues focus their means-end reasoning. In addition, it possible that, for experts, the overloaded plans do provide a savings in subsequent means-end reasoning, because experts may be more familiar with the infrequently used types of actions, and hence will not have as much difficulty as novices do in determining how to perform them. The experts are thus able reap the benefits of overloaded plans, without paying a heavy penalty in extra means-end reasoning, either when they initially identify the overloaded step or when they later elaborate it.

There are, of course, differences between the two empirical studies I have just discussed and the model of a practical-reasoning process being described in this paper. Hayes-Roth and Hayes-Roth were studying a situation in which an agent is provided with an overwhelming set of tasks all at once. Hayes studied just the opposite situation, in which an agent needed to construct a plan to achieve a single goal, that of machining a part. Moreover, in Hayes's study, though not in Hayes-Roth and Hayes-Roth's, the computation of a plan containing an overloaded intention presented a challenge, at least to novices. In contrast, I have been describing overloading as part of the ongoing practical-reasoning process in which agents continually engage as they encounter new goals in the presence of a rich structure of existing plans. I conjecture that, for many everyday practical-reasoning tasks, the recognition of overloading possibilities is not especially difficult; we are all, as it were, experts at quotidian reasoning. Finally, I am arguing in this paper not just that overloading occurs, but also that, given an agent's resource bounds, it is rational and it provides another reason, in addition to those proposed by Bratman (1987), for agents' committing themselves to the execution of certain plans in the first place.

4 The Rationality of Overloading

Let us consider in some more detail the claim that the process of overloading intentions, as set forth in the previous two sections, is a rational practical-reasoning strategy. First of all, it is important to be clear about just what the claim is: it is a claim about a reasoning strategy, and not about the outcome of each instance of that strategy. What I am claiming is that it is rational to use overloading as a way of generating solutions to means-end reasoning problems that are then fairly directly adopted, without being subject to further deliberation. I am *not* claiming that each plan thus adopted is optimal, in the sense of being the plan that would have been chosen had the agent engaged in complete practical reasoning, generating all possible plans and weighing them against one another. In other words, a plan adopted as a result of overloading may fail to be the one that would do the most to further the agent's goals, given her beliefs. To use the decision-theory slogan, it may fail to maximize her expected utility.

To see the problem, recall the way in which overloading works. Faced with a new goal, the agent examines her existing, adopted plans, attempting to find an action that can be overloaded; failing that she may attempt simple elaborations of her existing plans in an attempt to uncover an overloading possibility. If she succeeds in either of these

Situation	Employs Overloading	Adopted plan is Optimal	Complete Practical Reasoning *would have been* Worthwhile	Overloading "Pays"
1	N			–
2	Y	Y		Y
3	Y	N	N	Y
4	Y	N	Y	N

Table 1: Outcomes of Practical Reasoning

ways, finding a plan for her new goal that involves overloading, she directly adopts that plan (assuming its consistency with her already adopted plans), without generating and considering possible alternatives. But some of the alternatives she thereby overlooks might, in fact, have been judged better than the plan suggested by overloading. If so, the plan she adopted is, in an important sense, suboptimal.

Practical-reasoning "shortcuts", such as the use of overloading, seem inherently to raise the possibility of suboptimality. To judge the rationality of a practical-reasoning strategy, we thus need to ask whether, averaging over a sufficiently large number of uses of the strategy, the benefits derived from the avoidance of extra reasoning outweigh the disadvantages of any resulting suboptimality.

In (Bratman *et al.* 1988), we explored this question as it pertained to the strategy of "filtering," bypassing full deliberation concerning options that are believed to be incompatible with an agent's adopted plans. Our analysis centered on an enumeration of the possible outcomes of the use of filtering. I shall here present a similar analysis of overloading. Table 1 lists possible outcomes of practical-reasoning processes, with regard to the use of overloading.

In Situation 1, the agent is unable to make use of overloading in solving some means-end reasoning problem; this case thus need not concern us further here.

In Situation 2, the agent, using overloading, finds a solution to her means-end reasoning problem. In addition, in this case, the plan thus generated is optimal, in the following sense: were the agent to have generated all possible means to her current goal and then weighed them against one another using her deliberation mechanism, the plan that was generated using overloading would have been judged the best and, consequently, would have been adopted. This is the most favorable outcome: the agent has avoided unnecessary practical reasoning, while not falling prey to suboptimality.

Situations 3 and 4 are a bit more complex. Here the agent has found a solution to her means-end reasoning problem by using overloading, but it is not an optimal solution, as defined above. Even so, the overloading process may have been worthwhile. Recall the difficulty that agents in dynamic environments face, due to their resource bounds. The world may change while they are engaging in practical reasoning. Their reasoning thus has an associated cost, in that it precludes other useful activities that may no longer be possible by the time the reasoning is complete. We therefore need to weigh the savings in reasoning cost provided by overloading against the cost of the suboptimal plan relative to the optimal one. Situation 3 occurs when the former is greater than the latter: complete practical reasoning would not have been worthwhile in such cases. Thus, instances of Situation 3 are ones in which overloading turns out, on balance, to be beneficial. Situation 4 is the opposite case: there, overloading leads to a plan that is suboptimal, and the cost of the suboptimality is greater than the savings in reasoning cost. Thus, in Situation 4, overloading fails to pay: the agent would have been better, in such cases, to take the extra time to engage in complete practical reasoning.

One might be tempted to conclude that overloading is a rational process to the degree that it leads to instances of Situations 2 and 3, rather than Situation 4. This is not completely warranted, however. Instead, the analysis leads to a more modest claim, that overloading is rational, *relative to a strategy of complete deliberation*, to the degree that it leads to Situations 2 and 3, rather than Situation 4. But there may be other practical-reasoning strategies that have even better outcomes than overloading: if so, it may be more rational to use these alternatives. The challenge is to define such alternatives. Indeed, the optimal strategy may be one that combines overloading with other reasoning shortcuts, such as generating and directly adopting only the most commonly used plan for certain kinds of goals. Given both empirical evidence, such as that provided by Hayes's study, and commonsense intuition about the benefits of overloading, one quite plausible conjecture is that overloading will be part of any complete optimal strategy.[6]

Because putting an action to multiple uses tends to be a good thing to do, the likelihood of Situations 2 and 3 is increased by the very nature of overloading: the plan involving an overloaded action is likely to be, if not the very best alternative, at least a reasonably good one. Unfortunately, however, Situation 4 cannot, in general, be completely avoided. But an agent can minimize the frequency with which instances of Situation 4 occur by using overloading judiciously. For instance, an agent should not rely on overloading to bypass deliberation about plans to achieve goals that are highly significant. It would, for instance, be foolish for a high-school senior to adopt a plan to attend Harvard rather than any other college simply because she plans a sightseeing trip to the Boston area the following summer, and could overload her intention to travel there if she selects Harvard. Very important decisions should always be made on the basis of complete means-end reasoning and deliberation. Nonetheless, overloading can be quite useful for less critical, everyday practical-reasoning problems.

5 Intentions and Expectations

Up until now, I have depicted overloading as a process that exploits an agent's intentions. Yet, typically when an agent forms an intention to perform some action A, she also forms a belief, or expectation, that she will do A. My intention to swim today at lunchtime leads me to expect that I will swim today at lunchtime. Similarly, my intention to stop at Nan's office on the way to the pool generates an expectation that I will stop there.

Expectations are simply beliefs about future activities and circumstances. Whenever an agent engages in practical reasoning, she makes use of her expectations. Then, given the close association between intentions and expectations, we may be led to ask whether overloading really is a special form of practical reasoning, or whether it is just an example of the ordinary use of beliefs about the future to guide practical reasoning.

A partial answer to this question can be given if we consider whether an agent can have the intention to perform some action A without having the associated belief that she will in fact perform A. If this state of affairs is possible, and, moreover, the intention in question can be overloaded, then overloading is not simply a form of expectation use, as no expectation is involved here.

Although some philosophers, notably Grice (1971) and Harman (1986), have argued against the possibility of an agent's intending to A without believing she will do A, Bratman (1987), for one, disagrees. While accepting that an agent cannot intend to do A while believing that she will *not* do A, Bratman argues that the belief that the agent will do A can be missing. Here he has two kinds of cases in mind. First, there are cases in which an agent believes that she may fail to try to perform her intended action, as, for instance, with an agent who intends to stop at the bookstore on her way home from work, but believes that, because she is absent-minded, she is likely to forget to do so. Second, there are cases in which an agent believes that she will try to perform the intended action, but that when she does so, she is likely to fail. An example of this is an agent who intends to move an extremely heavy log, although doubting that she is, in fact, strong enough. Bratman marshals evidence in favor of viewing these as cases of intention: the agent, he argues, will filter options inconsistent with the performance of these actions, will engage in means-end reasoning about their achievement, and will track their success.

To the extent that one agrees with Bratman on this issue, one should also allow for the possibility of agents' overloading intentions that do not have associated expectations. Imagine Bratman's agent who intends to stop at the bookstore without believing that she will do so. This agent filters out options incompatible with stopping at the bookstore: she will not deliberate about taking an alternative route home, one that does not take her by the bookstore—and she will not form the intention to take such a route. She will engage in means-end reasoning related to her intended bookstore stop: she will plan to take along money to pay for her intended purchases. And so on. But if this agent is willing to allow her attitude toward stopping at the bookstore to shape her further practical reasoning in these ways, there seems to be no reason that she should not also exploit it in solving further means-end reasoning problems. If she should determine that she needs to buy some new disks for her home computer, it seems quite sensible for her to overload her intention to stop at the bookstore, using it in a plan to buy the disks.

Of course, in saying this, we must bear in mind that the agent can be no more confident of the success of an intention than she is of the correctness of each of the beliefs upon which that intention depends, as well as of the success of each of the intentions upon which it relies. As a result, a rational agent should not overload an intention whose success she doubts in elaborating a plan for an extremely important goal. Once again we see that overloading must be used judiciously: it may be primarily applicable to ordinary, everyday decision making.

So, if we accept that intentions can exist without associated expectations, we are justified in claiming that overloading is at least somewhat different from the ordinary use of beliefs about the future in practical reasoning. But this is a fairly weak conclusion, since most of the time our intentions probably are accompanied by associated expectations. There is, however, a stronger argument that can be made, namely, that there are advantages to overloading that do not arise in ordinary expectation use, and that these advantages lead to overloading being a different type of process than ordinary expectation use. To show this, we must first distinguish between three kinds of expectations that an agent may have.

First, there are those expectations that derive directly from an agent's intentions, expectations that she will do the actions she intends. Let us call these *primary expectations.*

Not all expectations are primary. Very often agents expect that the actions they intend will have side-effects—results that, while foreseen, are not intended. I expect that when I go swimming I will get chlorine in my hair, although I do not intend to get chlorine in my hair. I also expect that when I take my usual route to the pool, I will pass an automatic teller machine (ATM), although again, I do not intend to pass an ATM. One test that has been put forth as a way of distinguishing between intended effects and side-effects is that the latter are not "tracked": agents will not vary their actions to ensure the success of side-effects, at least in normal circumstances (Bratman 1987). Should I discover that the water in the pool in which I swim has been replaced with salt water, I will not attempt to locate a bottle of chlorine to pour into my hair. And should I find, while on my way to the pool, that the ATM has been moved, I will not alter my route in order to pass it. We can attach the label *secondary expectations* to beliefs that foreseen side-effects of intended actions will occur.

Secondary expectations, like primary ones, are rather closely related to an agent's intentions. Agents may also have expectations that are relatively divorced from any particular intentions they have. Here is an example: I expect that shortly before noon, Steve will pass by my office on his way to lunch. I do not intend to do any action of which this is either an intended effect or a side-effect. But experience has taught me to expect him to pass by. Expectations of this kind can be called *independent*. They do not directly concern the occurrence either of actions I intend or actions I expect to perform in doing what I intend.

There are thus three kinds of expectations that agents may have about future events. They expect to perform the actions they intend. They expect side-effects of the actions they intend to occur. And they expect many other, independent events to occur as well. In performing practical reasoning, agents may make use of all three kinds of expectations. But only primary and secondary expectations can be used in overloading. The use of independent expectations is merely what I have been calling ordinary expectation use, and differs qualitatively from overloading.

I have presented overloading as a process that uses intentions, but, at least for intentions that have associated primary expectations, nothing I have said would preclude viewing overloading as instead making use of the primary expectations. Let us then turn to secondary expectations. Suppose that I once again have a plan to go swimming at lunch, this time by myself, and that my usual route takes me past an ATM. Suppose, also, that I have once again forgotten to get money for lunch. Then my secondary expectation of passing the ATM does suggest a solution to my problem of getting money: I can stop and make a withdrawal at the ATM when I pass it. Moreover, it seems reasonable that I can once again presume that it is a good idea to make additional use of an action that I already expect to perform. Thus, I can directly adopt this plan to get money, without generating and considering alternatives. What was previously just a foreseen side-effect of an intended action thereby becomes a full-fledged intention. As a result, I will, for instance, track its success: should I now discover that the ATM has been moved, I may well alter my route to ensure that I pass it.

Note the difference between this case and the situation I described earlier, in which I overloaded my intention to meet Nan. There I had one intention—to meet Nan—which came to serve two distinct roles. I intended to meet her in order to go swimming with her, and I intended to meet her in order to get money for lunch. Meeting Nan is part of the elaboration of both of these goals. In the current example, my intention to pass the ATM is not overloaded. I intend it for a single purpose, in order to get money. But I expect that I would pass it anyway, even if I did not have the goal of getting money.

So in attempting to solve means-end reasoning problems, agents can examine their secondary expectations, about actions they will perform as side-effects of intended actions. When a (secondarily) expected action is found that can be put to use in an elaboration of the end in question, the agent can then form a new intention to perform that action as a way of achieving the end without further deliberation.

What about independent expectations? My expectation that Steve will pass by my office may be central in the plan I adopt to get money for lunch: I may decide to borrow it from him when he passes by. But does the formation of this plan really involve the overloading process as I have described it?

I claimed that when an agent exploits a secondary expectation, she forms, without complete deliberation, a new intention to perform the action that she had originally merely expected to perform as a side-effect of some other action. When I decide to make a withdrawal using the ATM, expecting that I will pass the ATM, I form the intention to pass the ATM. But when an agent forms a plan that relies on an independent expectation, she does not, in general, form an intention to bring about that expectation.[7] It seems unlikely, in the current example, that I would form an intention to ensure that Steve passes by my office shortly before noon.

But why not? Recall the interrelated reasons that overloading is worthwhile. First, plans suggested by overloading have a certain efficiency of action: they involve putting one action to more than one use. Second, the use of overloading generally leads to an efficiency of reasoning, as the agent can solve one means-end reasoning problem, about how to carry out the overloaded intention, in place of two (or more). These advantages together justify an agent's bypassing complete deliberation about the overloaded plan and alternatives to it, a move that results in further savings in reasoning cost. However, when an agent recognizes that an independent expectation can be useful in solving a means-end reasoning problem, the formation of a new intention to ensure that the expected event actually occurs would undercut the advantages of overloading. The newly formed intention would, at least initially, be in service of only one goal. And additional means-end reasoning would be required, to determine how to ensure the occurrence of the newly intended action, or at least, how to avoid unintentionally thwarting it. Thus, the agent is not, in this case, justified in bypassing deliberation. In contrast, when an agent forms an intention to perform a secondary expectation, at least the latter advantage still obtains, because she already intends to perform some action that will bring about the secondarily expected action; this then justifies direct adoption of the new intention.

There is, then, a fundamental difference between overloading and ordinary expectation use. The former offers certain advantages: if the agent is able to find a plan that involves overloading, she can reasonably presume that it will have an efficiency of action, and that it will cut down on the amount further means-end reasoning she will need to perform. She is thus justified in adopting it directly, thereby increasing even further the savings in reasoning cost. The same is not true of the ordinary use of independent expectations in solving practical-reasoning problems.

6 Implications for Plan Recognition

I have so far been focusing on overloading as a process that facilitates the generation and adoption of plans in resource-bounded agents.[8] But an understanding of the process can also be useful in other reasoning tasks that agents must perform, in particular, plan recognition.

Most agents do not live in isolation, but, rather, inhabit worlds in which they must interact with other agents. It is now generally accepted in AI that such interaction depends, in large part, on the abilities of agents to recognize one another's plans. As a result, the development of plan-recognition systems has been an active area of research (Allen 1983, Kautz 1990, Konolige and Pollack 1989, Pollack 1990). But most AI models of plan recognition have built upon traditional AI plan-generation techniques, in which an agent's plans are seen simply as recipes for achieving a goal, independent of and divorced from the role that plans play in the larger mental life of the agent. In particular, plan-recognition systems, like the classical plan-generation systems, have ignored the difficulties that dynamic environments pose for resource-bounded agents—and have ignored the ways in which planning can help agents cope with those difficulties. As a result, the capabilities of plan-recognition systems have been limited, a claim I will very briefly substantiate by means of two examples.

Here is the first: You speak with a colleague in the morning and determine that she intends to visit a nearby research laboratory later that day, and also plans to go to the bank sometime. You also know that she plans to attend the monthly departmental lunch at noon that day. At 11:45, you see her heading out to the parking lot. Because you have a paper that you want delivered to the research laboratory, you attempt to recognize her plan: is she on her way to the laboratory? If she is, you'll run after her to give her the paper. You know that it takes only about 10 minutes round-trip to get to the bank, but that when she goes to the laboratory, she'll probably be gone for more than an hour. In that case, your likely conclusion is that the plan she's currently executing is to go to the bank, since going to the laboratory now would be incompatible with her plan to attend the departmental lunch.

To draw this conclusion, a plan-recognition system would need to take into account the assumption that an agent's plans tend to be consistent with one another. And this assumption is justified by the claim that one way for a resource-bounded agent to curtail her reasoning is to consider only those options that are consistent with one's existing plans, a claim that arises, in turn, from consideration of the problems posed by resource boundedness (Bratman 1987, Bratman *et al.* 1988). Existing plan-recognition systems do not take this into account; indeed, they generally do not even consider the fuller set of adopted plans that the agent is known to have. Instead, they look only at the currently observed actions—in the current example, heading to the parking lot.

The first example thus illustrates that plan-recognition systems need to take into account characteristics of plans such as those described by Bratman (1987). A similar illustration can be given for the significance of the characteristics of plans that result from the overloading process. Imagine that late one afternoon, your husband calls you at work to offer to stop on the way home to pick up a video tape and to buy a pizza. After hanging up, you remember that you've invited some guests for dinner. You try to call back to ask him to buy a larger pizza, but he's already left. There are about a half dozen pizza parlours that you two typically frequent. However, one of these, Luigi's, is located in the same shopping center as the video-rental store. You thus conclude that your husband is probably going to stop at Luigi's, so you call there and have them tell him to buy a larger pizza.

As before, a plan-recognition system that considered only the intended action of buying a pizza would be unable to decide which pizza parlour your husband was intending to go to. To reason to that conclusion, a system would have to take into account his intention to go to the video-rental store, and the fact that agents tend to overload actions when possible. Again, the justification for overloading stems from a consideration of the challenge of resource boundedness, and the use of planning as a way of coping with that challenge.

In short, a better understanding of the roles that plans play in an agent's mental life, and of the constraints on plans that arise from those roles, can improve not only our models of plan generation, but also our models of plan recognition. The two examples just given suggest that, in performing plan recognition, an agent must take into account not just the observed actions, but also other actions and plans that she knows the observed agent intends. And to make use of this information, she must employ knowledge about the relations that exist among an agent's coexisting plans—relations that arise from the roles those plans play.

7 Conclusion

In his theory of intention, Bratman (1987) argued that it is rational for resource-bounded agents to form and adopt plans, because in so doing they lessen the amount of subsequent practical reasoning needed, and, hence, make it easier to cope with their dynamic environments. In this essay, I have tried to bolster that claim, by describing a practical-reasoning strategy I called overloading, in which an agent makes use of her existing plans to focus means-end reasoning and to bypass deliberation that would otherwise be necessary. I explored the rationale behind overloading, namely, the commonsense presumption that it is generally worthwhile to "kill two birds with one stone." And I claimed that an appreciation of the process of overloading, like other practical-reasoning strategies, can be beneficial to tasks such as plan recognition as well as plan generation.

There is also a more general demonstration made by this essay, namely, that AI models of intelligence can be usefully informed by philosophical theories—or, to put it the other way around, that (at least some) philosophical theories can be computationally realized. Bratman's theory of intention led to the high-level architectural specifications for an artificial agent given by Bratman, Israel, and me (1988), which in turn led to an implemented, prototype agent currently being used in an experimental study of tradeoffs in agent design (Pollack and Ringuette 1990). These three pieces of work also led to the development of the notion of overloading described in this paper. Overloading is essentially a computational process, one that a researcher designing an artificial agent could embody in algorithms incorporated in that agent. Philosophical theories can and do matter to the AI researcher.

Acknowledgments

This work was supported by the Office of Naval Research Center under Contract No. N00014-89-C-0095, by a contract with the Nippon Telegraph and Telephone Corporation, and by a gift from the System Development Foundation.

I am indebted to Michael Bratman for conversations that led to the development of the ideas presented here, and to him, Phil Cohen, Barbara Grosz, David Israel, and Kurt Konolige for valuable comments on an earlier draft of this paper.

Notes

[1] A similar claim might be made about the reasoning processes of biological agents, under the assumption that evolution plays the role of the designer.

[2] My use of the term *overloaded* derives from the technical sense it is given in computer science: one says that a symbol is overloaded when it is used for more than one purpose. For example, in most programming languages, the symbol "+" is overloaded to signify both integer and real addition. I do *not* mean to suggest that an overloaded intention is one that is, in some sense, too heavily burdened.

[3] So, to continue the recipe analogy: an agent may, by looking in a cookbook, find several different recipes for chocolate cake, without committing herself to make any of them, or even to make cake. And she may, by performing means-end reasoning, generate several plans for borrowing money, without intending one or another, and without, in fact, intending to borrow money.

[4] Early models of AI planning cast it as a search process in which the nodes represent states of the world, rather than partial plans (Fikes and Nilsson 1971). The model described in this paper is more general, subsuming the state-based approach. However, it is still a simplification. Techniques have been developed, especially within the past few years, for interleaving planning and execution, so that the select/elaborate cycle can be interrupted prior to completion, and the system can perform actions and engage in further reasoning under the influence of only partially specified plans (Georgeff 1987).

[5] As Hayes notes, operator overlap is closely related to Wilensky's (1983) notion of *partial plan overlap*.

[6] In fact, it may be possible to evaluate such conjectures experimentally. In some recent work, I have been developing a system for the experimental evaluation of alternative practical-reasoning strategies, called Tileworld (Pollack and Ringuette 1990). Tileworld consists of a simulated robot agent and a simulated environment. The experimenter can manipulate the characteristics of the environment, as well as the particular practical-reasoning strategies employed by the agent. Large numbers of simulations can then be automatically run, and the resulting performance analyzed to determine the advantages and disadvantages of particular reasoning strategies for environments exhibiting various features.

[7] Notice that it may often be impossible for an agent to form an intention whose object, strictly speaking, is the independent expectation. While secondary expectations have as their objects actions that the agent will herself perform, independent expectations are typically about circumstances in which she will find herself as a result of other agents' actions. I cannot form an intention whose object is the action of Steve's passing by my office shortly before noon; as Castaneda has argued, intentions are first-person attitudes (Castaneda 1975). This is why it is necessary to speak of agents' forming intentions to *bring about* the independent expectation.

[8] Much of this section derives from a talk that I gave at the Second Annual Workshop on Plan Recognition, American Association for Artificial Intelligence (AAAI), Detroit, MI, August 1989. I am grateful for feedback provided by the other workshop participants.

References

(Agre and Chapman 1987) Philip E. Agre and David Chapman. Pengi: An implementation of a theory of activity. In *AAAI-87, Proceedings of the Sixth National Conference on Artificial Intelligence*, Seattle, WA, 1987.

(Allen 1983) James F. Allen. Recognizing intentions from natural language utterances. In Michael Brady and Robert C. Berwick, editors, *Computational Models of Discourse*, pages 107–166. MIT Press, Cambridge, MA, 1983.

(Bratman *et al.* 1988) Michael E. Bratman, David J. Israel, and Martha E. Pollack. Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(4), 1988. Also to appear in John L. Pollock and Robert Cummins, eds., *Philosophy and AI: Essays at the Interface*, MIT Press, Cambridge, MA, 1991.

(Bratman 1987) Michael E. Bratman. *Intention, Plans and Practical Reason.* Harvard University Press, Cambridge, MA, 1987.

(Brooks 1991) Rodney A. Brooks. Intelligence without representation. To appear in *Artificial Intelligence*, 1991.

(Castaneda 1975) Hector-Neri Castaneda. *Thinking and Doing.* Reidel, Dordrecht, Holland, 1975.

(D'Ambrosio and Fehling 1989) Bruce D'Ambrosio and Michael Fehling. Resource-bounded agents in an uncertain world. In *Proceedings of the AAAI Symposium on Limited Rationality*, Stanford, CA, 1989.

(Doyle 1988) Jon Doyle. Artificial intelligence and rational self-government. Technical Report CS-88-124, Carnegie Mellon University, Pittsburgh, PA, 1988.

(Fikes and Nilsson 1971) Richard E. Fikes and Nils J. Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208, 1971.

(Georgeff 1987) Michael P. Georgeff. Planning. In Joseph Traub, editor, *Annual Review of Computer Science, Vol. 2*. Annual Reviews, Inc., Palo Alto, CA, 1987.

(Grice 1971) H.P. Grice. Intention and uncertainty. *Proceedings of the British Academy*, 57:263–179, 1971.

(Harman 1986) Gilbert Harman. *Change in View.* MIT Press, Cambridge, MA, 1986.

(Hayes-Roth and Hayes-Roth 1979) Barbara Hayes-Roth and Frederick Hayes-Roth. A cognitive model of planning. *Cognitive Science*, 3:275-310, 1979.

(Hayes 1989) Caroline Hayes. A model of planning for plan efficiency: Taking advantage of operator overlap. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, MI, 1989.

(Jeffrey 1983) Richard Jeffrey. *The Logic of Decision, 2nd edition.* University of Chicago Press, Chicago, 1983.

(Kaelbling 1988) Leslie P. Kaelbling. Goals as parallel program specifications. In *AAAI-88, Proceedings of the Seventh National Conference on Artificial Intelligence*, St. Paul, MI, 1988.

(Kautz 1990) Henry A. Kautz. A circumscriptive theory of plan recognition. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, editors, *Intentions in Communication.* MIT Press, Cambridge, MA, 1990.

(Konolige and Pollack 1989) Kurt Konolige and Martha E. Pollack. Ascribing plans to agents: Preliminary report. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Detroit, MI, 1989.

(Pollack and Ringuette 1990) Martha E. Pollack and Marc Ringuette. Introducing the Tileworld: Experimentally evaluating agent architectures. In *AAAI-90, Proceedings of the Ninth International Conference on Artificial Intelligence*, Boston, MA, 1990.

(Pollack 1990) Martha E. Pollack. Plans as complex mental attitudes. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack, editors, *Intentions in Communication*. MIT Press, Cambridge, MA, 1990.

(Pollock 1989) John L. Pollock. Oscar: A general theory of rationality. In *Proceedings of the AAAI Symposium on Limited Rationality*, Stanford, CA, 1989.

(Sacerdoti 1977) Earl D. Sacerdoti. *A Structure for Plans and Behavior*. American Elsevier, New York, 1977.

(Savage 1972) Leonard J. Savage. *The Foundations of Statistics, 2nd edition*. Dover Press, New York, 1972.

(Tate 1977) Austin Tate. Generating project networks. In *Proceedings of IJCAI-77*, Cambridge, MA, 1977.

(Wilensky 1983) Robert Wilensky. *Planning and Understanding*. Addison-Wesley, Reading, MA, 1983.

(Wilkins 1988) David E. Wilkins. *Practical Planning: Extending the Classical AI Paradigm*. Morgan Kaufmann, San Mateo, CA, 1988.